A Supervised Learning Algorithm for Binary Domain Classification of Web Queries using SERPs

Alexander C. Nwala and Michael L. Nelson
Old Dominion University, Department of Computer Science
Norfolk, Virginia, 23529
{anwala, mln}@cs.odu.edu

arXiv:1605.00184v2 [cs.IR] 3 May 2016

ABSTRACT

General purpose Search Engines (SEs) crawl all domains (e.g., Sports, News, Entertainment) of the Web, but sometimes the informational need of a query is restricted to a particular domain (e.g., Medical). We leverage the work of SEs as part of our effort to route domain specific queries to local Digital Libraries (DLs). SEs are often used even if they are not the "best" source for certain types of queries. Rather than tell users to "use this DL for this kind of query", we intend to automatically detect when a query could be better served by a local DL (such as a private, access-controlled DL that is not crawlable via SEs). This is not an easy task because Web queries are short, ambiguous, and there is lack of quality labeled training data (or it is expensive to create). To detect queries that should be routed to local, specialized DLs, we first send the queries to Google and then examine the features in the resulting Search Engine Result Pages (SERPs), and then classify the query as belonging to either the ***scholar*** or ***non-scholar*** domain. Using 400,000 AOL queries for the ***non-scholar*** domain and 400,000 queries from the NASA Technical Report Server (NTRS) for the ***scholar*** domain, our classifier achieved a precision of 0.809 and F-measure of 0.805.

Keywords

Search Engines, Digital Libraries, Web queries, Query Understanding, Web query domain classification.

1. INTRODUCTION

In this paper we focus on domain classification of queries which targets two classes - the ***scholar*** domain and the ***non-scholar*** domain. The ***scholar*** domain targets queries associated with academic or research content. For example, queries such as "*fluid dynamics*", "*stem cell*", "*parallel computing*" belong to the ***scholar*** domain, while queries such as "*where to find good pizza*", "*bicycle deals*", and "*current weather*" belong to the ***non-scholar*** domain. SEs are often used even if they are not the "best" source for certain types of queries. For example, according to the Ithaka S+R Faculty Survey 2015 [18], faculty members are likely to begin their research discovery with a general purpose SE. In fact, according to this survey, over 30% of the surveyed faculty report SEs as the start point of their academic research discovery [18]. Rather than tell users to "use this DL for this kind of query", we intend to automatically detect when a query could be better served by a local DL (such as a private, access-controlled DL that is not crawlable via SEs). We propose a novel method which does not rely on processing the actual query. Instead, we trained a classifier based on the features found in a Google SERP (Search Engine Result Page). The classifier was trained and evaluated (through 10-fold cross validation) on a dataset of 600,000 SERPs evenly split across both classes. The results were validated by 200,000 SERPs evenly split across both classes yielding a classification precision of 0.809 and F-measure of 0.805. We targeted a binary class, however our method could be scaled to accommodate other classes if the right features are found.

The rest of the paper is organized as follows: in Section 2, we present research similar to ours but outline how our work differs from the state of the art. Section 3 gives an overview of our solution in two stages: building the classifier and classifying a query. In Section 4, we identify the features and discuss the rationale for selecting the features. Section 5 outlines the classifier building/evaluation and classification tasks. We present our results in Section 6, discuss our findings in Section 7, and finally in Section 8 we conclude.

2. RELATED WORK

The problems of query domain classification and query routing to specialized collections have been studied extensively. Jingbo et al. [19] built a domain knowledge base from web pages for query classification. Gravano, et al. [3] built a classifier targeting the geographical locality domain. Dou Shen et al. [15] built an ensemble-search based approach for query classification in which queries are enriched with information derived from Search Engines. Wang et al. [17] in an effort to label queries, investigated the use of semi-supervised learning algorithms built from HTML lists to automatically generate semantic-class lexicons. For the query understanding task (which includes domain classification), Liu et al. [9] proposed a method which takes advantage of the linguistic information encoded in hierarchical parse trees for query understanding. This was done by extracting a set of syntactic structural features and semantic dependency features from query parse trees. Jansen et al. [6] presented a classification of user intent. Three hierarchical intent classes were considered: informational, navigational, and transactional intent.

Sugiura and Etzioni [16] introduced Q-pilot: a system which dynamically routes user queries to their respective specialized search engines. Similarly, in order to improve search results, Lee et al. [8] proposed two types of features (user-click behavior and anchor-link distribution) in order to discover the "goal" behind a user's Web query. Jian Hu

et al. [5] proposed a method of discovering a user query's intent with Wikipedia in order to help a Search Engine to automatically route the query to some appropriate Vertical Search Engines. In this work, the Wikipedia concepts were used as the intent representation space (each intent domain is represented as a set of Wikipedia articles and categories). Thereafter, the intent of incoming queries were identified through mapping the query into the Wikipedia representation space.

The query domain classification problem is not new, and discovering the domain of a query facilitates routing the query to an appropriate Search Engine or Digital Library. However, since queries are short, ambiguous, and in constant flux, maintaining a labeled training dataset is expensive. Similar to other methods, our method utilized the Google Search Engine, but differs from the state of the art by not processing the query directly, but the SERP instead. We used Google SERPs to develop our method. However, our techniques could be applied to other Search Engines such as Bing, once the corresponding features extracted from Google SERPs are identified on the Bing SERPs. Fig. 8 identifies some features on the Bing SERP which we collect from Google SERPs.

3. METHOD OVERVIEW OF LEARNING ALGORITHM FOR DOMAIN CLASSIFICATION

Our solution to the web query binary domain classification problem can be summarized in two stages:

1. **Stage 1 - Building the classifier:**

   First, identify the discriminative features. Second, build a dataset for the ***scholar*** domain class and ***non-scholar*** domain class. Third, train a classifier. Fourth, evaluate the classifier using 10-fold cross validation.

2. **Stage 2 - Classifying a query:**

   First, issue the query to Google and download the SERP. Second, extract the features. Finally, use the classifier hypothesis function to make a prediction.

4. EXTRACTING FEATURES FROM SERPS:

After extensive study (information theory and visualization), we identified 10 features to be extracted from the Google SERP:

1. **Knowledge Entity (f_1):**

   Domain: Binary, $f_1 \in \{0, 1\}$

   Figure Example when Present: Fig. 3

   Figure Example when Absent: Fig. 1

   Description: f_1 indicates the presence or absence of the Google Knowledge Graph Panel. The presence of this panel (Fig. 3) often indicates the presence of a real world object such as a person, place, etc. as understood by Google.

2. **Images (f_2):**

   Domain: Binary, $f_2 \in \{0, 1\}$

   Figure Example when Present: Fig. 2

   Figure Example when Absent: Fig. 1

   Description: f_2 indicates the presence or absence of Images on the SERP.

3. **Google Scholar (f_3):**

   Domain: Binary, $f_3 \in \{0, 1\}$

   Figure Example when Present: Fig. 2

   Figure Example when Absent: Fig. 1

   Description: f_3 indicates the presence or absence of a citation. A citation of a research paper or book is often a strong indicator about a web query of the ***scholar*** class (Table 2).

4. **Ad ratio (f_4):**

   Domain: $\mathbb{R}$, $f_4 \in [0, 1]$

   Figure Example when Present: Fig. 1

   Figure Example when Absent: Fig. 2

   Description: f_4 represents the proportion of ads on the SERP. Even though Search Engines strive to include ads (in order to pay their bills), our information analysis showed that the ***non-scholar*** class SERPs had the higher proportion of ads (Fig. 4).

5. ***non-HTML* rate (f_5):**

   Domain: $\mathbb{R}$, $f_5 \in [0, 1]$

   Figure Example when Present: Fig. 1 and Fig. 2

   Figure Example when Absent: NA

   Description: f_5 represents proportion of *non-html* types on the SERP as reported by Google (as opposed to dereferencing the links and retrieving the "Content-Type" header values). We defined *non-html* types to include members of set **M**.

   $\mathbf{M} = \{$*pdf, ppt, pptx, doc, docx, txt, dot, dox, dotx, rtf, pps, dotm, pdfx*$\}$

   For example, if a given SERP has results with the following Mime types:

   ```
   1.  html
   2.  html
   3.  html
   4.  pdf  (non-html)
   5.  html
   6.  ppt  (non-html)
   7.  html
   8.  pptx (non-html)
   9.  html
   10. html
   ```

   Since there are three *non-html* types, $f_5 = \frac{3}{10}$

   Our intuition which expects the ***scholar*** class to have more *non-html* types was matched by Fig. 5.

6. **Vertical permutation (f_6):**

   Domain: $\mathbb{Z}$, $f_6 \in [1, 336]$

   Figure Example when Present: Fig. 1 and Fig. 2

   Figure Example when Absent: NA

   Description: f_6 represents the rank from ${}^8P_3$ (8 Permutation 3) possible page order of the SERP. In addition to the default SERP (All), the Google SERP provides a set of eight custom vertical search pages:

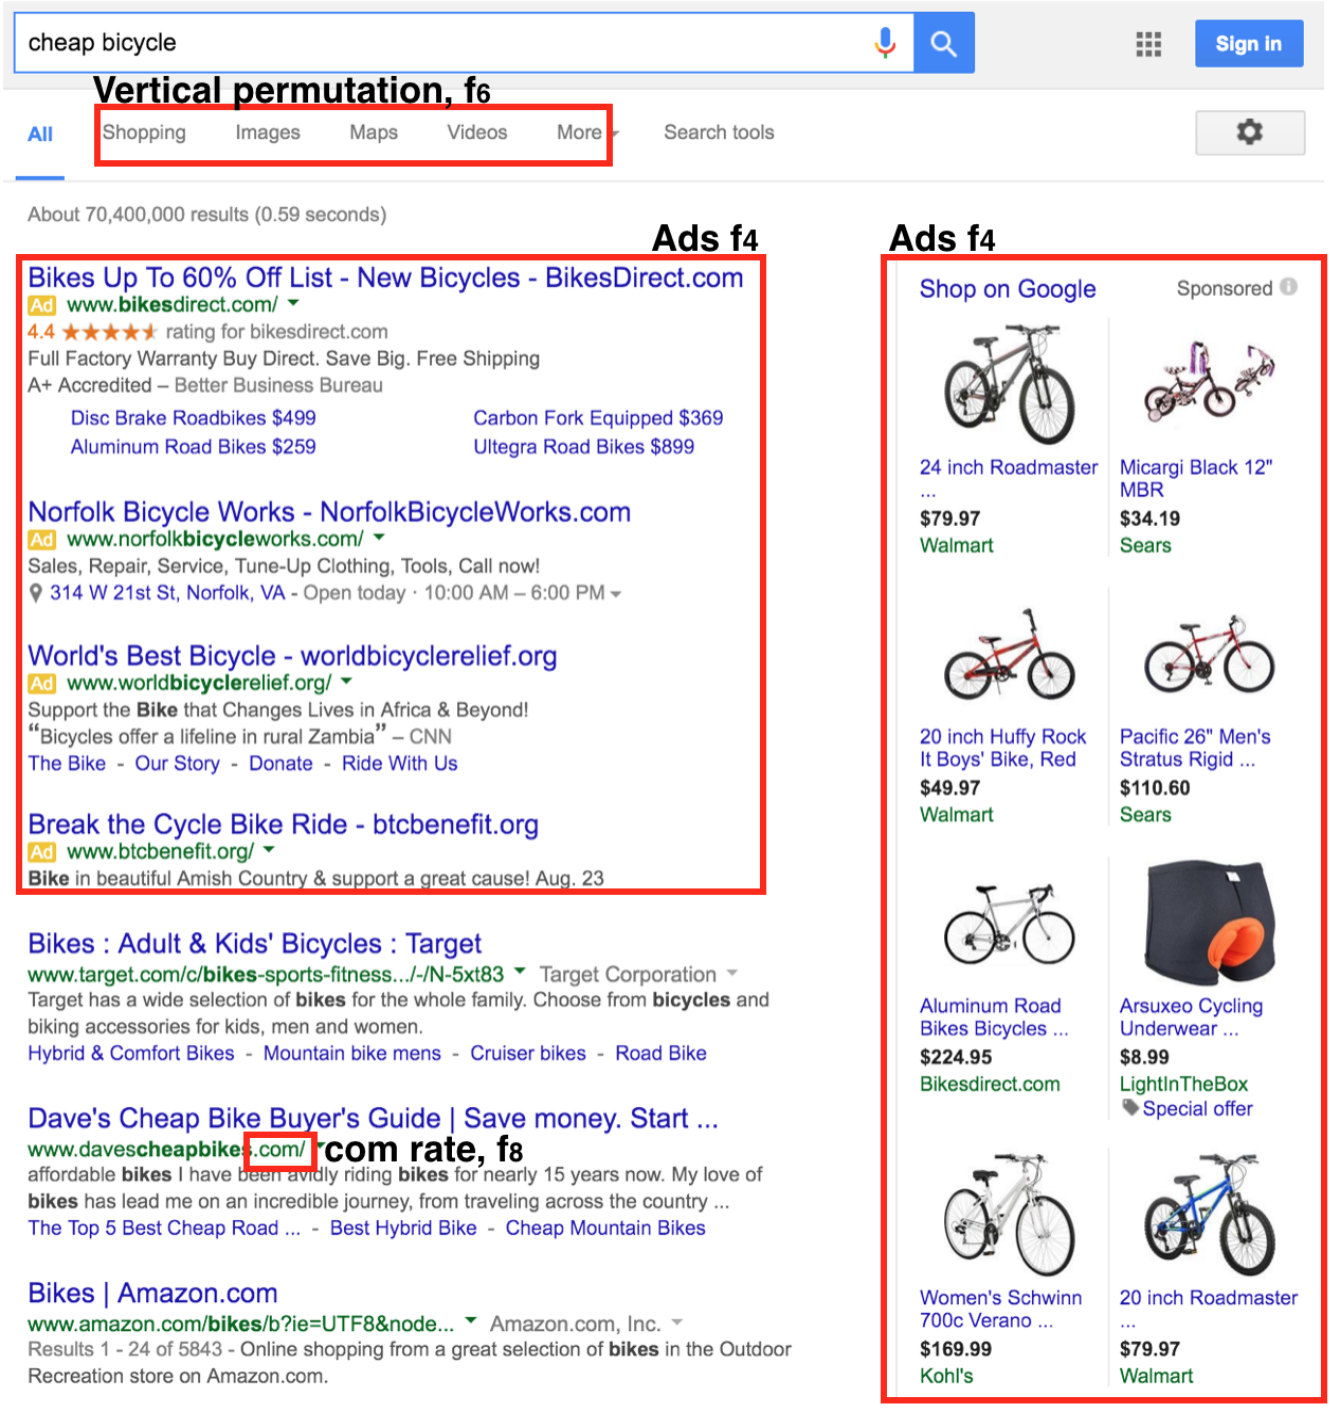

Figure 1: ***non-scholar*** queries typically feature "shopping" in the Vertical permutation (f_6), Ads (f_4), and are dominated by *com* TLDs (f_8). Snapshot date: March 15, 2016.

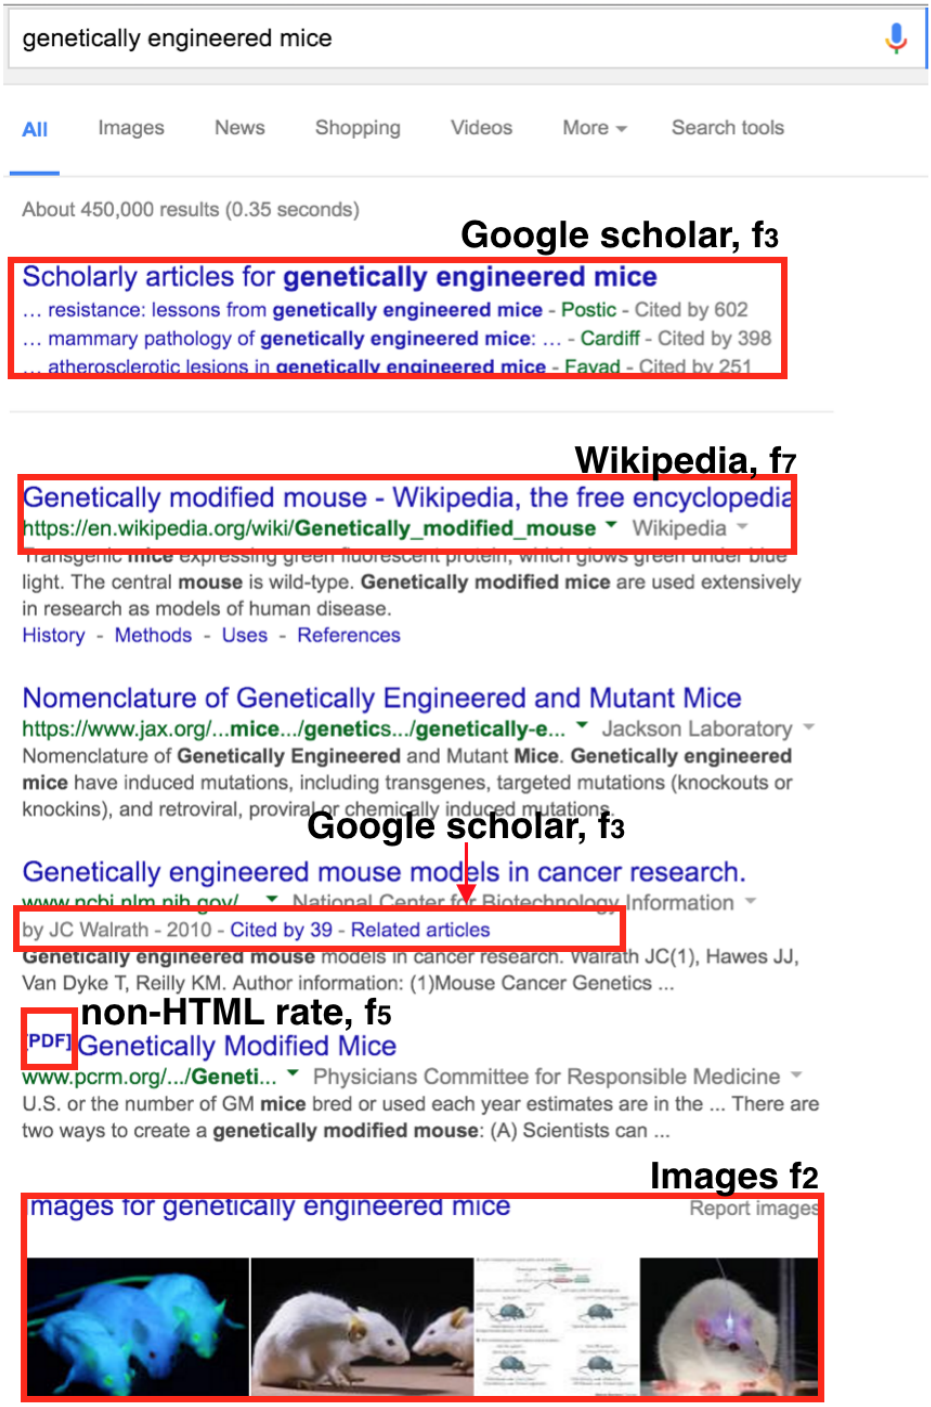

Figure 2: ***scholar*** queries typically feature results from Wikipedia (f_7) and Google Scholar (f_3), as well as *non-html* types (f_5). Note: Google feedback panel was removed from image to accommodate more features. Snapshot date: March 15, 2016.

Table 1: Complete list of features

Feature	**Domain**	**Information Gain**	**Gain Ratio**
f_1, Knowledge Entity	{0, 1}	0.013	0.033
f_2, Images	{0, 1}	0.021	0.024
f_3, Google Scholar	{0, 1}	0.248	0.298
f_4, Ad ratio	$\mathbb{R} \in [0, 1]$	0.051	0.022
f_5, *non-HTML* rate	$\mathbb{R} \in [0, 1]$	0.205	0.048
f_6, Vertical permutation	$\mathbb{Z} \in [1, 336]$	0.100	0.020
f_7, Wikipedia	{0, 1}	0.033	0.034
f_8, *com* rate	$\mathbb{R} \in [0, 1]$	0.127	0.092
f_9, Max. Title Dissimilarity	$\mathbb{R} \in [0, 1]$	0.018	0.007
f_{10}, Max. Title Overlap	$\mathbb{Z} \in [0, \max_{t_i} \lvert q' \cap t_i \rvert]$	0.031	0.013

Table 2: Proportion of SERPs with a Knowledge Entity (f_1), Images (f_2), Google Scholar Citation (f_3), and a Wikipedia link (f_7), for ***scholar*** and ***non-scholar*** classes for 200,000 SERPs evenly split across both classes.

Feature Title	**Scholar**	**non-Scholar**
Knowledge Entity (f_1)	4.10%	10.97%
Images (f_2)	38.27%	22.60%
Google Scholar (f_3)	50.51%	2.54%
Wikipedia (f_7)	49.63%	28.72%

Table 3: Summary Statistics for Ad ratio (f_4), *non-HTML* rate (f_5), Vertical permutation (f_6), *com* rate (f_8), Max. Title Dissimilarity (f_9) and Max. Title Overlap (f_{10}), for ***scholar*** and ***non-scholar*** classes for 200,000 SERPs evenly split across both classes.

	Scholar		**non-Scholar**	
Feature Title	**Mean**	**Std. Dev**	**Mean**	**Std. Dev**
Ad ratio (f_4)	0.0470	0.1100	0.1060	0.1680
non-HTML rate (f_5)	0.1660	0.1910	0.0170	0.0630
Vertical permutation (f_6)	228.9220	59.8090	240.3640	56.4760
com rate (f_8)	0.2740	0.2230	0.1000	0.1790
Max. Title Dissimilarity (f_9)	0.8420	0.0870	0.8250	0.0870
Max. Title Overlap (f_{10})	1.8570	1.1660	1.8810	1.4630

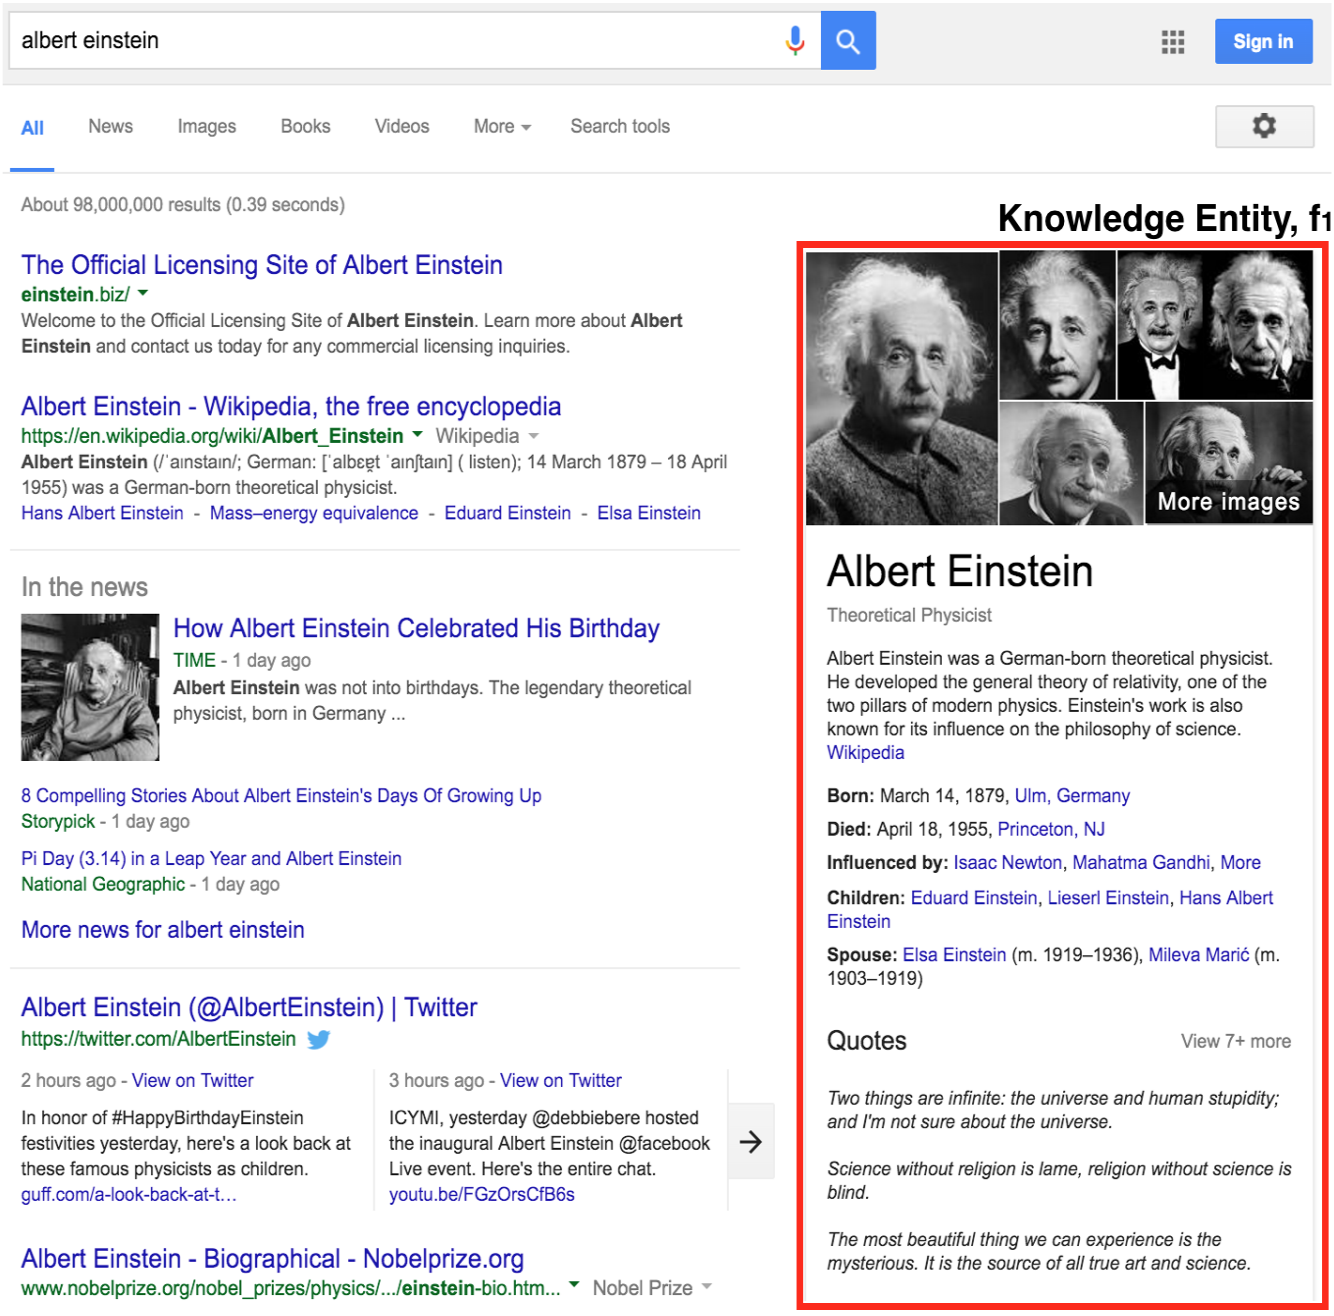

Figure 3: Google Knowledge Graph Panel (f_1) for "albert einstein" query. Snapshot date: March 15, 2016.

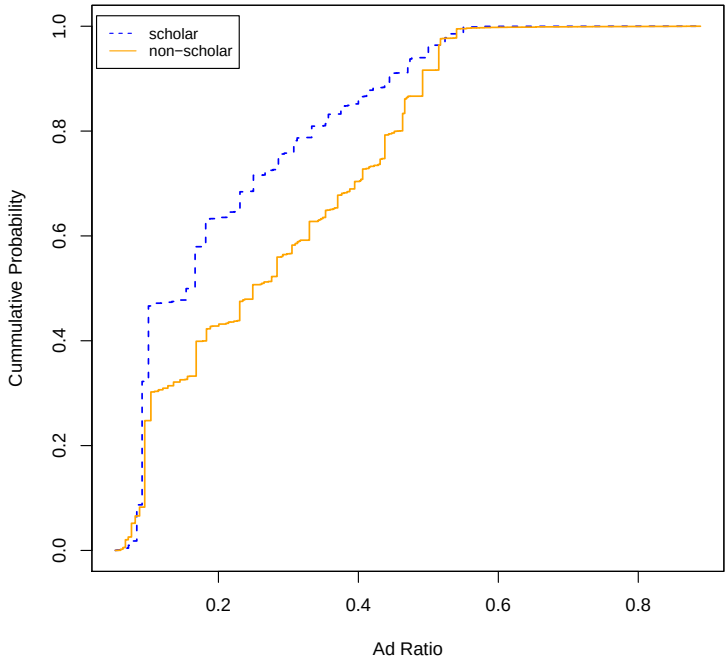

Figure 4: CDF of Ad Ratio (f_4) for ***scholar*** (blue, dashed) & ***non-scholar*** (orange, solid) classes for 200,000 SERPs evenly split across both classes; ***non-scholar*** with a higher Probability of ads.

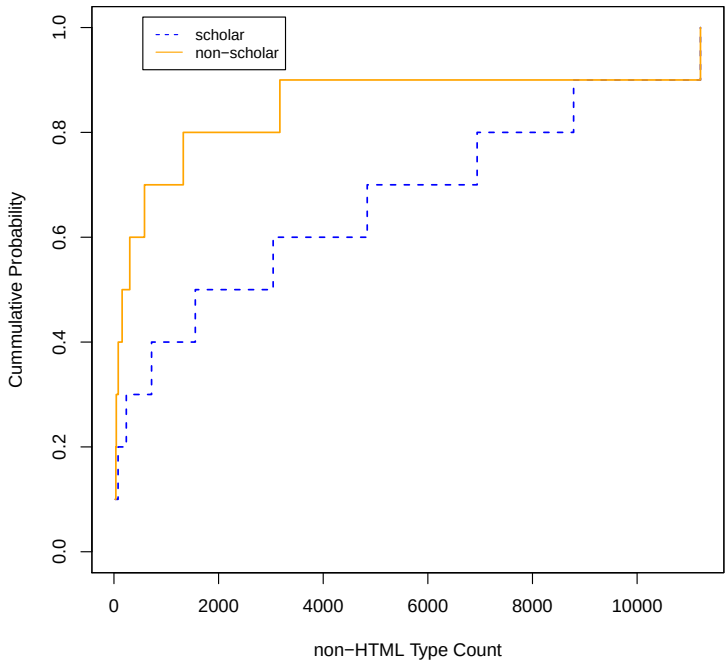

Figure 5: CDF of *non-html* types (f_5), ***scholar*** (blue, dashed) and ***non-scholar*** (orange, solid) classes for 200,000 SERPs evenly split across both classes; ***scholar*** with a higher probability of *non-html* types.

`Apps, Books, Flights, Images, Maps, News, Shopping, and Videos`

Even though there are eight possible vertical pages, only four are displayed simultaneously with the default SERP (All). The rest are hidden. Different queries trigger different permutations of four from the set of eight. The order in which the four pages are displayed is tied to the kind of query.

For example, the query "*cheap bicycle*" (Fig. 1) presents:

`Shopping, Images, Maps, and Videos`

And the query "*genetically engineered mice*" (Fig. 2) presents:

`Images, News, Shopping, and Videos`

In order to capture this information we assigned each page a number based on its position in alphabetical order:

```
0 - Apps, 1 - Books, 2 - Flights, 3 - Images,
4 - Maps, 5 - News, 6 - Shopping, 7 - Videos
```

We empirically determined that the relevance of a query to the vertical pages decreases as we go from the first vertical page to the last vertical page. For example, the query "*cheap bicycle*" is most relevant to the *Shopping* vertical page, and less relevant to the *Images* vertical page and so on. Due to the decrease in relevance we decided to capture only the first three pages. This gives a total of 336 (${}^8P_3$) permutations (in lexicographic order) beginning from Rank 0 to Rank 335:

```
Permutation 0:   0 1 2: Apps, Books, Flights
Permutation 1:   0 1 3: Apps, Books, Images
Permutation 2:   0 1 4: Apps, Books, Maps
.
.
.
Permutation 335: 7 6 5: Videos, Shopping, News
```

7. **Wikipedia (f_7):**

   Domain: Binary, $f_7 \in \{0, 1\}$

   Figure Example when Present: Fig. 2

   Figure Example when Absent: Fig. 1

   Description: f_7 indicates the presence or absence of a Wikipedia link on the SERP. Based on our study, the ***scholar*** class presented a higher proportion of SERPs with a Wikipedia link (Table 2).

8. ***com* rate (f_8):**

   Domain: $\mathbb{R}$, $f_8 \in [0, 1]$

   Figure Example when Present: Fig. 1 and Fig. 2

   Figure Example when Absent: NA

   Description: f_8 represents the proportion of *com* links on the SERP. The motivation for this feature was due to observations from visualizing our datasets which revealed that queries of the ***scholar*** class had a lesser proportion of *com* links on the SERP, instead had a higher proportion of TLDs such as *edu* and *gov*.

9. **Maximum Title Dissimilarity (f_9):**

   Domain: $\mathbb{R}$, $f_9 \in [0, 1]$

   Figure Example when Present: Fig. 1 and Fig. 2

   Figure Example when Absent: NA

   Description: f_9 represents the maximum dissimilarity value between the query and all SERP titles: Given a query q, with SERP title t_i, longest SERP title T, Levinshtein Distance function LD, and cardinality (length) of longest SERP title $|T|$, then:

$$f_9 = \max_{t_i} \frac{LD(q, t_i)}{|T|} \tag{1}$$

   For example, considering the query "cheap bicycle" in Fig. 1, the maximum f_9 (dissimilarity) score is 0.8182, which belongs to the second title - "Dave's Cheap Bike Buyer's Guide | Save money. Start ...". For the query "genetically engineered mice" (Fig. 2), the maximum f_9 is 0.6557 which belongs to the first title - "Genetically modified mouse - Wikipedia, the free encyclopedia". See Section 5.3 for an example of how f_9 was calculated for a given query.

10. **Maximum Title Overlap (f_{10}):**

    Domain: $\mathbb{Z}$, $f_{10} \in \mathbb{Z} \in [0, \max_{t'_i} |q' \cap t'_i|]$

    Figure Example when Present: Fig. 1 and Fig. 2

    Figure Example when Absent: NA

    Description: In contrast with f_9 which captures dissimilirity, f_{10} represents the cardinality of the maximum common value between the query and all SERP titles: Given a query set q' with SERP title set t'_i, then:

$$f_{10} = \max_{t'_i} |q' \cap t'_i| \tag{2}$$

    For example, for the query "cheap bicycle" (Fig. 1), the maximum f_{10} score is 1 which belongs to the second title - "Dave's Cheap Bike Buyer's Guide | Save money. Start ...". For the query "genetically engineered mice" (Fig. 2), the maximum f_{10} score is 3 which belongs to the second title - "Nomenclature of Genetically Engineered and Mutant Mice". See Section 5.3 for an example of how f_{10} was calculated for a given query.

Individually, many features may not be helpful in distinguishing the ***scholar*** query from the ***non-sholar*** query, however it is in the combination of these features in which we could learn to classify web queries based on the combined signals these features indicate.

Figures 1 and 2 present a subset of the features for both classes while Table 1 presents the complete list of features. Note the ads (f_4) present in the ***non-scholar*** SERP and the absence of PDF (f_5) documents (Fig. 1). For the ***scholar*** SERP (Fig. 2), note the presence of a Wikipedia link (f_7), the PDF (f_5) document and the Google Scholar article (f_3). Therefore, at scale, we could learn from not just the presence of a feature, but its absence.

5. CLASSIFIER TRAINING AND EVALUATION:

After feature identification, we downloaded the Google SERPs for 400,000 AOL 2006 queries [14] and 400,000 NTRS (from 1995-1998) queries [12] for the non-Scholar and Scholar datasets, respectively. We thereafter extracted the features described in Section 4, and then trained a classifier on the dataset. After exploring multiple statistical machine learning methods such as SVMs, we chose logistic regression because of its fast convergence rate and simplicity.

5.1 Logistic Regression:

Logistic regression is analogous to linear regression. Linear regression describes the relationship between a continuous response variable and a set of predictor variables, but logistic regression approximates the relationship between a discrete response variable and a set of predictor variables. The description of logistic regression in this section is based on Larose's explanation [7].

5.1.1 Logistic Regression Hypothesis Representation:

Given a predictor variable X, and a response variable Y, the logistic regression hypothesis function (outputs estimates of the class probability), is the conditional mean of Y given $X = x$, is denoted as $P(Y|x)$. The hypothesis function represents the expected value of the response variable for a given value of the predictor and is outlined by Eqn. 3

$$P(Y|x) = \frac{e^{\beta_0+\beta_1 x}}{1+e^{\beta_0+\beta_1 x}}; 0 \leq P(Y|x) \leq 1 \tag{3}$$

$P(Y|x)$ is nonlinear and sigmoidal (S-shaped), which is unlike the linear regression model which assumes the relationship between the predictor and response variable is linear. Since $0 \leq P(Y|x) \leq 1$, $P(Y|x)$ may be interpreted as a probability. For example, given a two class problem where Y corresponds to the presence of a disease ($Y = 1$) or the absence of the same disease ($Y = 0$),

$P(Y = 1|x) + P(Y = 0|x) = 1$. Thus $P(Y|x)$ may be threshold at 0.5:

If $P(Y|x) \geq 0.5$; Y = 1

If $P(Y|x) < 0.5$; Y = 0

5.1.2 Logistic Regression Coefficients Estimation - Maximum Likelihood Estimation:

The least-squares method offers a means of finding a closed-form solution for the optimal values of the regression coefficients for linear regression. But this is not the case for logistic regression - there is no such closed-form solution for estimating logistic regression coefficients. Therefore, we may apply the *maximum likelihood estimation* method to find estimates of the parameters β_i which maximize the likelihood of observing the data. Consider the likelihood function $l(\beta|x)$ (Eqn. 4),

Let $P(Y|x) = \alpha(x)$

$$l(\beta|x) = \prod_{i=1}^{n} [\alpha(x_i)]^{y_i} [1-\alpha(x_i)]^{1-y_i} \tag{4}$$

which represents the parameters β_i that describe the probability of the observed data, x. To discover the maximum likelihood estimators, we must find the values of the parameters β_i $(\beta_0, \beta_1, \beta_2, ..., \beta_m)$, which maximize the likelihood function $l(\beta|x)$.

The log likelihood of $l(\beta|x)$, $L(\beta|x)$ is often preferred to the likelihood function, because it is more computationally tractable, thus expressed by Eqn. 5:

$$L(\beta|x) = ln[l(\beta|x)] = \sum_{i=1}^{n} \{y_i ln[\alpha(x_i)] + (1-y_i) ln[1-\alpha(x_i)]\} \tag{5}$$

To find the *maximum likelihood estimators* we have to differentiate $L(\beta|x)$ with respect to each parameter. Thereafter, we set the result equal to zero and solve the equation. For linear regression a closed-form solution for the differentiation exists, but this is not the case for logistic regression, so me must resort to other numerical methods such as *iterative weighted least squares* [10]

5.2 Weka Logistic Regression:

Using Weka [4], we built a Logistic Regression Model (Eqns. 6, 7) on a 600,000 dataset evenly split across both classes. The model was evaluated using 10-fold cross validation yielding a classification precision of 0.809 and F-Measure of 0.805

$$p_q = \frac{e^{g_q}}{1+e^{g_q}} \tag{6}$$

$$g_q = 2.7585 + \sum_{i=1}^{10} c_i f_i \tag{7}$$

$c_i \in C$: The coefficient of feature f_i.

C: The coefficients matrix output by Weka contains the coefficients c_i for each feature f_i.

$$C = \begin{bmatrix} c_1 \\ c_2 \\ c_3 \\ c_4 \\ c_5 \\ c_6 \\ c_7 \\ c_8 \\ c_9 \\ c_{10} \end{bmatrix} = \begin{bmatrix} 0.8266 \\ -1.1664 \\ -2.7413 \\ -1.7444 \\ 6.2504 \\ -0.0017 \\ -1.0145 \\ -1.5367 \\ 1.8977 \\ -0.1737 \end{bmatrix}$$

To find the domain $d \in$ {Scholar, non-Scholar} of a query q, download Google SERP, then apply Algorithm 1.

Algorithm 1 Classify Query

Input: Google SERP s and query, q

Output: *scholar* or *non-scholar* domain label for q

function ClassifyQuery(s, q) ▷ SERP s for query q

1. f_1 ... $f_{10} \leftarrow s$ ▷ Initialize features
2. $g_q = 2.7585 + \sum_{i=1}^{10} c_i f_i$ ▷ Eqn 3; estimate g_q
3. $p_q = \frac{e^{g_q}}{1+e^{g_q}}$ ▷ Eqn 4; estimate class probability p_q

if $p_q \geq 0.5$ **then** ▷ Set domain

4. **return** *scholar domain*

else

5. **return** *non-scholar domain*

5.3 Classifying a Query (Example):

Table 4: f_9 and f_{10} calculated for SERP Titles for Query "*moon shot*". Note that titles t_1 and t_2 are similar but not exact.

t_i	**Title**	f_9	f_{10}
t_1	Moon Shot - Wikipedia, the free encyclopedia	0.7955	2
t_2	Moonshot - Wikipedia, the free encyclopedia	0.8372	0
t_3	HP Moonshot System | HP Official Site	0.7568	0
t_4	What is moonshot? - Definition from WhatIs.com	0.8478	0
t_5	Thinking big: Larry Page talks moon shots - Google Careers	0.8448	1
t_6	Moon Shots Program	0.5000	1
t_7	Moon Shot - Amazon.com	0.5909	2
t_8	Moonshot!, John Sculley - Amazon.com	0.8056	0
t_9	Google's Larry Page on Why Moon Shots Matter	0.7955	1
t_{10}	Google's Moon Shot - The New Yorker	0.7429	2
t_{11}	Moon Shots for Management	0.6400	1
	Maximum	0.8478	2

Consider the process of classifying the query: "*moon shot*" based on Algorithm 1. This example references Fig. 7. Also for this example, when a binary feature is present, we represent the feature value with 0 and when it is absent, we represent the feature value with 1. This is done in order to conform to Weka's convention, such that the result is comparable with Weka's result.

1. **Issue the query to Google and download SERP (Fig. 7)**
2. **Initialize all feature values f_1 ... f_{10}:**
   (a) Knowledge Entity, $f_1 = 1$ (Absent)
   (b) Images, $f_2 = 0$ (Present)
   (c) Google Scholar, $f_3 = 1$ (Absent)
   (d) Ad ratio, $f_4 = 0.0833$ (1 ad out of 12 results)
   (e) *non-HTML* rate, $f_5 = 0$ (SERP types has no member in set M (Section 4))
   (f) Vertical permutation, $f_6 = 275$ (*Shopping - 6, Images - 3, Videos - 7*)
   (g) Wikipedia, $f_7 = 0$ (Present)
   (h) *com* rate, $f_8 = 0.6363$ (7 *com* Tlds out of 11 links)
   (i) Maximum Title Dissimilarity, $f_9 = 0.8478$ (See Table 4.):
   Given query q = "moon shot",
   Fourth SERP title, t_4 = "What is moonshot? - Definition from WhatIs.com",
   Maximum SERP title length, $|T| = |t_4| = 46$,
   Levinshtein Distance between q and t_4, $LD(q, t_4) = 39$
   $f_9 = \max_{t_i} \frac{LD(q,t_4)}{|T|} = \frac{39}{46} = 0.8478$

Table 7: Logistic Regression Model Confusion Matrix

	scholar	**non-scholar**
scholar	TP: 224,360	FP: 41,266
non-scholar	FN: 75,640	TN: 258,734

Table 8: Validation: Logistic Regression Model Confusion Matrix

	scholar	**non-scholar**
scholar	TP: 74,651	FP: 13,473
non-scholar	FN: 25,349	TN: 86,527

   (j) Maximum Title Overlap, $f_{10} = 2$ (See Table 4.):
   Given query q = "moon shot", transformed to query set, q' = {"moon", "shot"},
   SERP title t_1 = "Moon Shot - Wikipedia, the free encyclopedia", transformed to SERP title set,
   t'_1 = {"shot", "wikipedia,", "-", "free", "moon", "encyclopedia", "the"},
   $f_{10} = \max_{t'_i} |q' \cap t'_i|$
   = {"moon", "shot"} ∩ {"shot", "wikipedia,", "-", "free", "moon", "encyclopedia", "the"} = 2

3. **Use Eqn 3. to estimate g_q:**
   $g_q = 2.7585 + (0.8266.f_1) + (-1.1664.f_2) + (-2.7413.f_3) + (-1.7444.f_4) + (6.2504.f_5) + (-0.0017.f_6) + (-1.0145.f_7) + (-1.5367.f_8) + (1.8977.f_9) + (-0.1737.f_{10})$
   $= 2.7585 + (0.8266 \times 1) + (-1.1664 \times 0) + (-2.7413 \times 1) + (-1.7444 \times 0.0833) + (6.2504 \times 0) + (-0.0017 \times 275) + (-1.0145 \times 0) + (-1.5367 \times 0.6363) + (1.8977 \times 0.8478) + (-0.1737 \times 2)$
   $= 0.5147$
4. **Use the Logistic Regression Hypothesis (Eqn 4.) to estimate the class probability p_q:**
   $p_q = \frac{e^{g_q}}{1+e^{g_q}} = \frac{e^{0.5147}}{1+e^{0.5147}} = 0.6259$
5. **Classify query:**
   Since $p_q \geq 0.5$, domain = ***scholar***

6. RESULTS

The logistic regression model was built using a 600,000 dataset split evenly across the ***scholar*** and ***non-scholar*** classes. The stratified 10-fold cross validation method was used to evaluate the model. The model produced a precision (weighted average over classes) of 0.809 and F-measure of 0.805 (Table 5). The model produced the confusion matrix outlined in Table 7.

In order to verify if the model was in fact as good as the preliminary results showed, we validated the model with a fresh "unseen" (not used to train/test the classifier) dataset of 200,000 evenly split across both the ***scholar*** and ***non-scholar*** class. The validation model yielded a precision of 0.810 (weighted average over classes) and F-measure of 0.805 (Table 6, Fig. 6). The model produced the confusion matrix outlined in Table 8.

Table 5: Detailed Accuracy by Class for Logistic Regression Model

Class	TP Rate	FP Rate	Precision	Recall	F-measure	ROC Area	
scholar	0.748	0.138	0.845	0.748	0.793	0.890	
non-scholar	0.862	0.252	0.774	0.862	0.816	0.890	
	0.805	**0.195**	**0.809**	**0.805**	**0.805**	**0.890**	Weighted Avg.

Table 6: Detailed Accuracy by Class for Validation Logistic Regression Model

Class	TP Rate	FP Rate	Precision	Recall	F-measure	ROC Area	
scholar	0.747	0.135	0.847	0.747	0.794	0.891	
non-scholar	0.865	0.253	0.773	0.865	0.817	0.891	
	0.806	**0.194**	**0.810**	**0.806**	**0.805**	**0.891**	Weighted Avg.

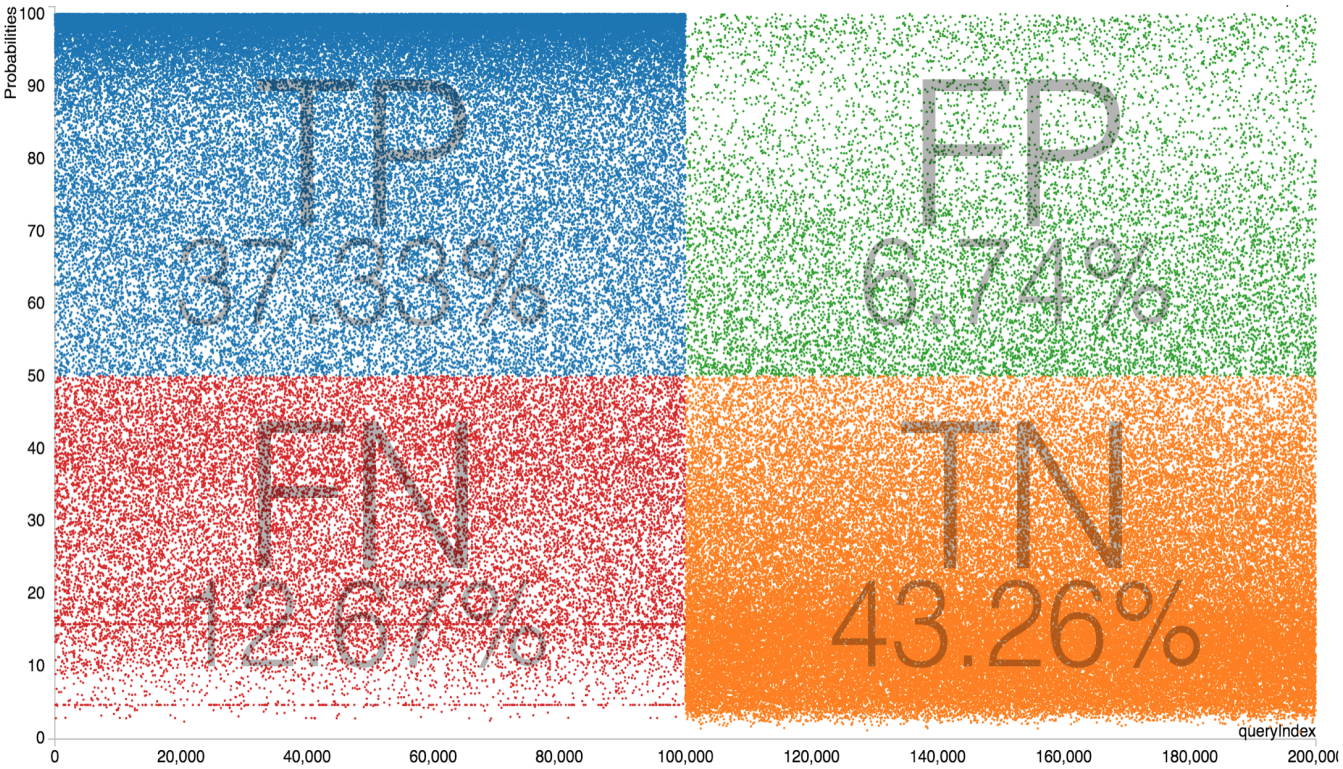

Figure 6: Distribution of Validation Model Class Probabilities for 200,000 Queries evenly split across ***scholar*** and ***non-scholar*** classes: Blue (Row 1, Col. 1, TP) ***scholar***, Orange (Row 2, Col. 2, TN) ***non-scholar***, Green (Row 1, Col. 2, FP), Red (Row 2, Col. 1, FN).

Table 9: Some Impure Queries in NTRS (***scholar***) Dataset and AOL (***non-scholar***) Dataset which Contribute to Classification Errors

	Scholar		**non-Scholar**	
Index	**Query**	**Class Probability**	**Query**	**Class Probability**
1	team and manage	0.1325	team and manage	0.1325
2	source supplies	0.1351	source supplies	0.1351
3	solour	0.1661	solour	0.1661
4	Pinkerton	0.1914	Pinkerton	0.1914
5	prying	0.2075	prying	0.2075
6	screen saver	0.2726	screen saver	0.2726
7	race car tests	0.2830	race car tests	0.2830
8	guaranteed funding	0.2888	guaranteed funding	0.2888
9	fix	0.4927	fix	0.4927

moon shot

Web Shopping Images Videos News More Search tools

About 25,100,000 results (0.32 seconds)

HP Moonshot System - Speed, Scale, and Specializationâ€Ž
Ad www.hp.com/â€Ž
4.5 rating for hp.com
Learn More About HP **Moonshot** Today.

Moon Shot - Wikipedia, the free encyclopedia
https://en.wikipedia.org/wiki/**Moon_Shot** Wikipedia
Moon Shot: The Inside Story of America's Race to the Moon is a book written by Mercury Seven astronaut Alan Shepard, with NBC News correspondent Jay ...

Moonshot - Wikipedia, the free encyclopedia
https://en.wikipedia.org/wiki/**Moonshot** Wikipedia
Moon shot or **moonshot** may refer to: The United States Apollo program; The abortive Soviet **Moonshot**; **Moonshot** (Google), a Google project to be realized ...

HP Moonshot System | HPÂ® Official Site
www8.hp.com/us/en/products/servers/**moonshot**/ Hewlettâ€'Packard
HP **Moonshot** helps you get more out of your infrastructure with less space, power consumption, and complexity â€" giving you the freedom to innovate and stay ...

What is moonshot? - Definition from WhatIs.com
whatis.techtarget.com › Topics › Computer Science › Robotics
A **moonshot**, in a metaphorical sense, is an ambitious, exploratory and ground-breaking project undertaken without any expectation of near-term profitability or ...

Images for moon shot Report images

More images for moon shot

Thinking big: Larry Page talks moon shots - Google Careers
https://www.google.com/about/.../thinking-big-larry-page.html Google
Whether it's creating self-driving cars or translating the web into 64 languages, we love big bets. Our company culture encourages experimentation and the free ...

Moon Shots Program
www.cancer**moonshots**.org/
Which cancer will you end? The University of Texas MD Anderson Cancer Center has a bold plan to quickly improve survival rates for many of the deadliest ...

See results about

Moon Shot: The Inside Story of America's Race to the ...
Published: 1994
Authors: Jay Barbree, Deke Slayton, Alan Shepard

Feedback

Figure 7: Portion of SERP for "moon shot" query, used for the example in Section 5.3

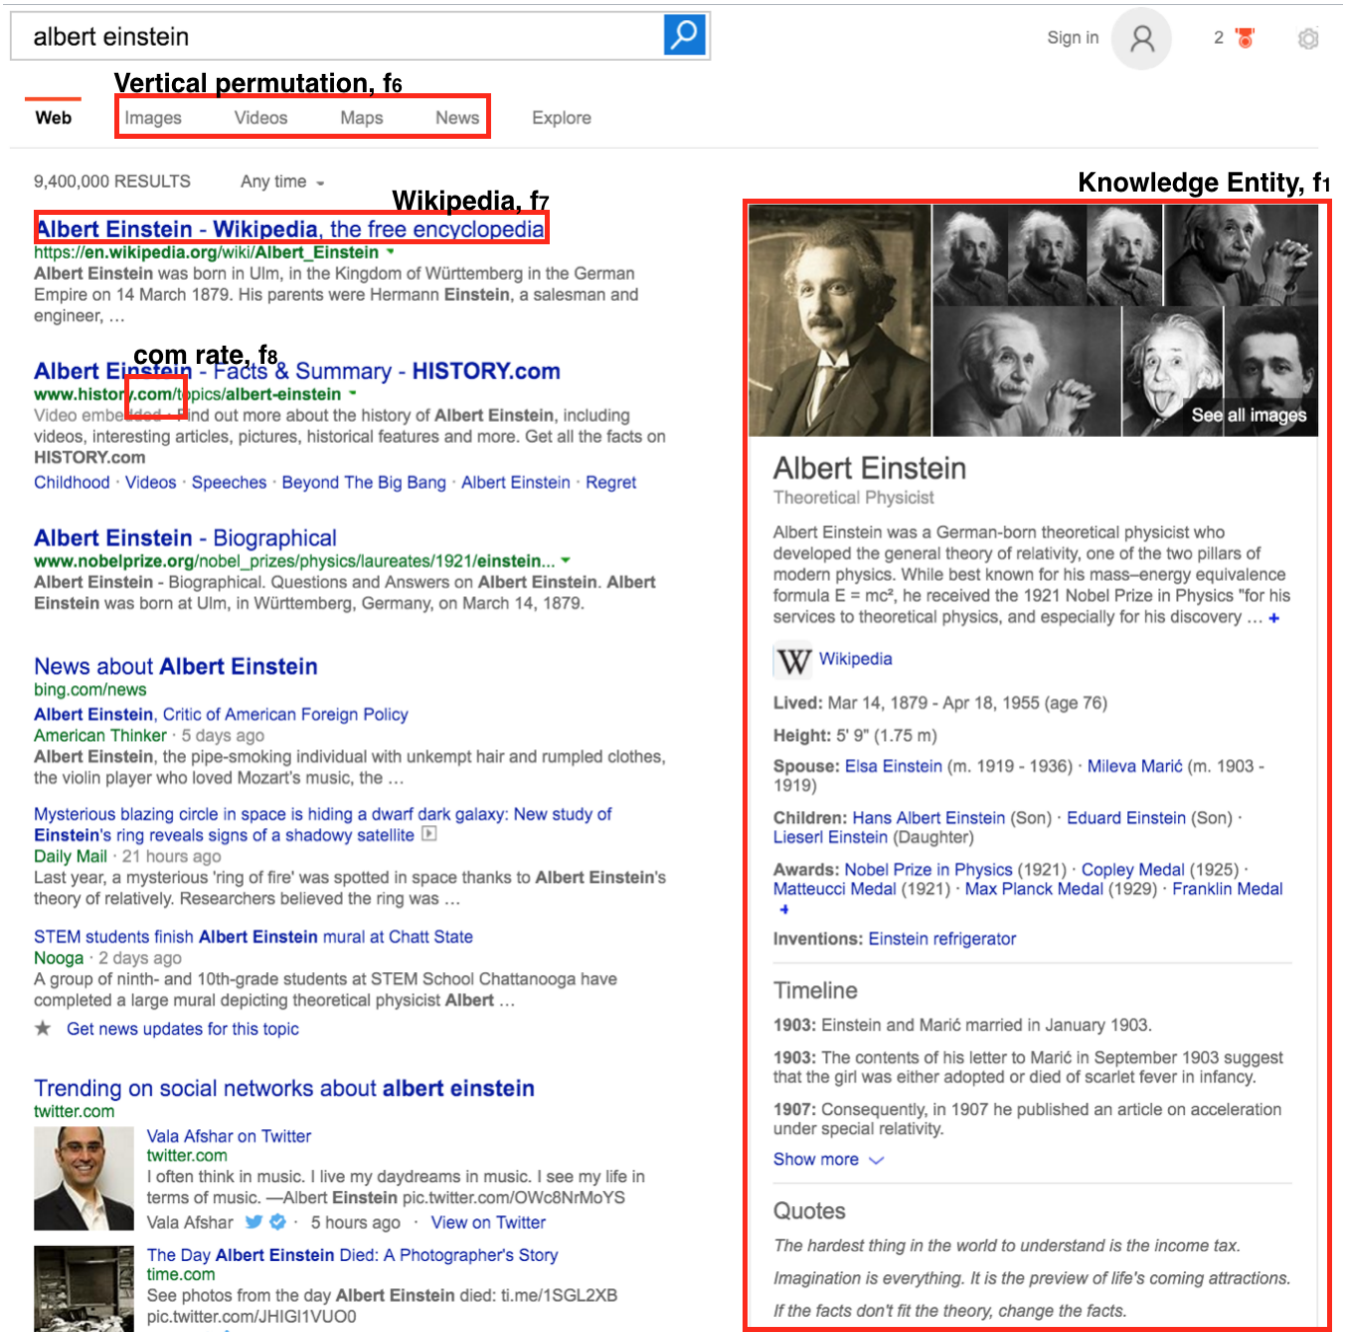

Figure 8: Bing SERP with for "albert einstein" query with features f_1, f_6, f_7 and f_8. Snapshot date: April 16, 2016.

7. DISCUSSION

Google is essentially a black box. Callan [2] explored how Search Engines search a number of databases, and merge results returned by different databases by ranking them based on their ability to satisfy the query. Similarly, Murdock and Lalmas [11] explored aggregated search, addressing how to select, rank and present resources in a manner that information can be found efficiently and effectively. Arguello et al. [1] addressed the problem of vertical selection which involves predicting relevant verticals for user queries issued to the Search Engine's main web search page. Even though these methods ([2], [11], and [1]) may shed some light about some operations associated with Search Engines, we still do not know about all the mechanism of the Google search, but we can see what it returns. Also, instead of replicating the methods discussed in Section 2 and methods such as [2], [11], and [1], we leverage the results produced by Search Engines.

Before building our dataset, we had to be able to detect and extract the features from SERPs. This was achieved by scraping the Google SERP. While identifying the features was an easy task, the limitation of scraping is that if the HTML elements (e.g., tags, attributes) of the SERP change, the implementation might need to change too. However, based on empirical evidence, SERPs tend to change gradually. But if they do, for example, if Google adds a new Vertical or a new source of evidence, this could result in a new feature for our method, therefore our implementation [13], has to be updated. Our method processes Google SERPs, but it could easily be adapted to other SERPs such as Bing. Fig. 8 outlines some features from the Google SERP that carry over to Bing.

Before building our model, we visualized the datasets at our disposal and performed information analysis (information gain/gain ratio - Table 1) in order identify features which could contribute discrimination information for our classifier. After identifying 10 features (Table 1), we trained and evaluated a logistic regression model on the dataset and achieved a precision of 0.809 and F-measure of 0.805. Our method is not without limitations. For example, the datasets are presumed "pure." But this is not the case since Scholar queries exist in the non-Scholar dataset and vice versa, thus contribute to classification errors. However, it should be noted that our model performed very well even though it was trained on an "impure" dataset (e.g., Table 9).

The NTRS dataset (***scholar*** class) had queries which could be easily identified as ***non-scholar*** queries. For example, consider Table 9 which presents some queries found in the NTRS dataset with questionable membership, thus contributing to classification errors. The low class probabilities could signal that the classifier had little "doubt" when it placed the ***non-scholar*** label (correct based on human judgement) on such queries, but was penalized due to the presence of these queries in the ***scholar*** dataset. Therefore, our Classifier could be improved if the "impure" queries and their respective SERPs are removed.

Similar to the NTRS dataset, the AOL 2006 dataset (***non-scholar*** class) also had queries which could be easily identified as ***scholar*** queries. For example, Table 9 presents some queries found in the AOL dataset with questionable membership, thus contributing to classification errors. It is clear from the high class probabilities, that the classifier had little "doubt" when it placed the ***scholar*** label (correct based on human judgement) on such queries, but was penalized due to the presence of these queries in the ***non-scholar*** dataset. The python open source code for our implementation and the NTRS dataset is available for download at [13].

8. CONCLUSIONS

In order to route a query to its appropriate digital library, we established query understanding, (thus domain classification) as a means toward deciding if a query is better served by a given digital library. The problem of domain classification is not new, however the state of the art consists of a spectrum of solutions which process the query unlike our method which leverages quality SERPs. We define a set of ten features in SERPs that indicate if the domain of a query is scholarly or not. Our logistic regression classifier was trained and evaluated (10-fold cross validation) with a SERP corpus of 600,000 evenly split across the ***scholar***, ***non-scholar*** classes, yielding a precision of 0.809 and F-measure of 0.805. The performance of the classifier was further validated with a SERP corpus of 200,000 evenly split across both classes yielding approximately the same precision and F-measure as before. Even though our solution to the domain classification problem targets two domains (***scholar*** and ***non-scholar***), our method can be easily extended to accommodate more domains once discriminative and informative features are identified.

9. REFERENCES

[1] J. Arguello, F. Diaz, J. Callan, and J.-F. Crespo. Sources of evidence for vertical selection. In *Proceedings of the 32nd international ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 315–322. ACM, 2009.

[2] J. Callan. Distributed information retrieval. In *Advances in Information Retrieval*, pages 127–150. Springer, 2002.

[3] L. Gravano, V. Hatzivassiloglou, and R. Lichtenstein. Categorizing web queries according to geographical locality. In *Proceedings of the Twelfth International Conference on Information and Knowledge Management*, pages 325–333. ACM, 2003.

[4] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The WEKA data mining software: an update. *ACM SIGKDD Explorations Newsletter*, 11(1):10–18, 2009.

[5] J. Hu, G. Wang, F. Lochovsky, J.-t. Sun, and Z. Chen. Understanding user's query intent with wikipedia. In *Proceedings of the 18th International Conference on World Wide Web*, pages 471–480. ACM, 2009.

[6] B. J. Jansen, D. L. Booth, and A. Spink. Determining the informational, navigational, and transactional intent of web queries. *Information Processing & Management*, 44(3):1251–1266, 2008.

[7] D. T. Larose. *Data mining methods & models*. John Wiley & Sons, 2006.

[8] U. Lee, Z. Liu, and J. Cho. Automatic identification of user goals in web search. In *Proceedings of the 14th International Conference on World Wide Web*, pages 391–400. ACM, 2005.

[9] J. Liu, P. Pasupat, Y. Wang, S. Cyphers, and J. Glass. Query understanding enhanced by hierarchical parsing structures. In *Automatic Speech Recognition and*

Understanding (ASRU), 2013 IEEE Workshop on, pages 72–77. IEEE, 2013.

[10] P. McCullagh and J. A. Nelder. *Generalized linear models*, volume 37. CRC press, 1989.

[11] V. Murdock and M. Lalmas. Workshop on aggregated search. In *ACM SIGIR Forum*, volume 42, pages 80–83. ACM, 2008.

[12] M. L. Nelson, G. L. Gottlich, D. J. Bianco, S. S. Paulson, R. L. Binkley, Y. D. Kellogg, C. J. Beaumont, R. B. Schmunk, M. J. Kurtz, A. Accomazzi, and O. Syed. The NASA technical report server. *Internet Research*, 5(2):25–36, 1995.

[13] A. Nwala. Query Classification Source Code. https://github.com/oduwsdl/QueryClassification, 2016.

[14] G. Pass, A. Chowdhury, and C. Torgeson. A picture of search. In *Proceedings of the 1st International Conference on Scalable Information Systems*, 2006.

[15] D. Shen, R. Pan, J.-T. Sun, J. J. Pan, K. Wu, J. Yin, and Q. Yang. Q2C@UST: our winning solution to query classification in KDDCUP 2005. *ACM SIGKDD Explorations Newsletter*, 7(2):100–110, 2005.

[16] A. Sugiura and O. Etzioni. Query routing for web search engines: Architecture and experiments. *Computer Networks*, 33(1):417–429, 2000.

[17] Y.-Y. Wang, R. Hoffmann, X. Li, and J. Szymanski. Semi-supervised learning of semantic classes for query understanding: from the web and for the web. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management*, pages 37–46. ACM, 2009.

[18] C. Wolff, A. B. Rod, and R. C. Schonfeld. Ithaka S+R US Faculty Survey 2015. http://www.sr.ithaka.org/wp-content/uploads/2016/03/SR_Report_US_Faculty_Survey_2015040416.pdf, 2016.

[19] J. Yu and N. Ye. Automatic web query classification using large unlabeled web pages. In *Web-Age Information Management, 2008. WAIM'08. The Ninth International Conference on*, pages 211–215. IEEE, 2008.